REPORT DOCUMENTATION PAGE

OMB No. 0704-0188

6890325

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	10/2/03	Final 1/10/00 – 1/9/03

4. TITLE AND SUBTITLE

Spike-Based Hybrid Computers

5. FUNDING NUMBERS

N00014-00-1-0244

6. AUTHOR(S)

Prof. Rahul Sarpeshkar

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Research Laboratory of Electronics
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research
Ballston Centre Tower One
800 North Quincy St.
Arlington, VA 22217-5660

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

00PRO4219-00

11. SUPPLEMENTARY NOTES

The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

14. SUBJECT TERMS	15. NUMBER OF PAGES
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Final Report

Office of Naval Research Grant No. N00014-00-1-0244

Grant Title: Spike-Based Hybrid Computers

PI Name: Rahul Sarpeshkar

Address: Massachusetts Institute of Technology Room 38-294, 77 Massachusetts Avenue Cambridge, MA 02139

Phone: (617) 258-6599
Fax: (617) 253-2510

E-mail Address: rahuls@mit.edu

Web Site: http://avnsl.mit.edu

Program Officer: Thomas McKenna

Concise Progress Summary:

This report documents progress made during the period July 1st 2002 to Jan 2003 for the ONR grant "Spike-Based Hybrid Computers". It is the fourth and final progress report for this grant. Our progress this year was in three areas 1) Ultra-Low-Power Time-Based Analog-to-Digital Conversion inspired by neuronal signal representations 2) An analog integrated circuit memory element for use in learning and adaptive systems, and 3) A predictive comparator with adaptive control that exploits knowledge of the signal to predict it and thus save energy.

Our novel time-based analog-to-digital converter uses a successive sub-ranging technique such that the conversion time increases linearly with the number of bits of precision. In contrast, in the traditional time-based converter of electronics, the dual-slope converter, the conversion time increases exponentially with the precision of the converter. In our prototype design of last year, the converter consisted of 3 matched capacitors, 2 unmatched reference currents, a single comparator with finite pulse-width output, a state-machine, and n one-bit flip-flops. A 1.5um AMI MOSIS chip consumed an area of 4.84 square mm, yielded 7 to 8 bits of precision with a conversion frequency ranging from 31.25 kHz to 62.5 kHz, and consumed 200 uW of power. Over the past year, we've been able to simplify the algorithm and completely eliminate a capacitor and reference current. As a result, we've reduced the overall power consumption and improved the conversion time and precision. Our latest 0.35 um TSM MOSIS chip consumed an area of 8.4 square mm, yielded 10 bits of precision with a conversion frequency up to 100 kHz, and consumed 150 uW. We are currently awaiting a 0.18um TSM MOSIS chip that explores slight modifications to the algorithm in addition to circuit topologies that reduce the noise and jitter. The latter chip is designed for 2 Volt operation, further reducing the overall power by at least 30 percent. We expect the final version of our converter, that we plan to publish and patent this year, to be the most power-efficient converter ever reported in the speed and precision range that we are exploring.

Our work on analog storage over the past year has been focused on developing an ultra-low power medium-term analog storage cell. There are two reasons why we

have concentrated on medium-term storage over long-term storage. The first reason for this focus is that our past research on capacitive-refresh storage has not conclusively shown that long-term analog storage techniques will offer a power or area savings over traditional digital methods. However, it has shown that simple leakage reduction or cancellation techniques can offer huge wins in terms of power and area over digital techniques in situations where long-term analog storage is not necessary (i.e. in circuits that adapt and learn from their environment in real time). The second reason for this focus is that an ultra-low power medium-term storage cell can be used as the basic storage element in a fully-refreshed long-term storage cell, offering an energy savings to the overall system due to its lower refresh rate. The primary result of the past year¡¦s work has been the development of an ultra-low power medium-term analog storage cell that is on the brink of becoming the world's leading non-floating gate implementation. Our cell can operate on less than 50nW with a 500us sampling time, and has been shown to achieve 11-bits of sampled precision. We expect to achieve 12-bits of sampled precision in our final version with the removal of a non-fundamental constant error. The cell operates on a 3.3V rail and can store voltages anywhere between 0.2V up to 3.1V, giving it a 2.9V input range. Additionally, during the storage phase, the cell's stored value remains within 12-bits of accuracy of the initial sampled value for over 45 seconds. This represents an increase of a factor of three in hold-time over the leading competitor in the literature. Additionally, the power consumption of this cell is over three orders of magnitude less than the leading competitor. After the 12-bit threshold of initial accuracy is broken, this cell will be the world's leading implementation of medium-term analog storage in all respects. This memory element will prove to be an invaluable tool in developing adaptive and learning neural systems in VLSI.

We added a linear predictor and adaptive control to a comparator to greatly reduce its delay as we described last year: In an experimental 1.5um VLSI chip implementation, our scheme enabled an almost 100x improvement in the power-delay product of a simple comparator. This year, we determined that the difference equations that govern our adaptive scheme may be represented by a unimodal discrete map. Therefore, near the limits of the scheme's stable regime of operation, we were able to experimentally confirm that there was a period-doubling route to chaos. Our scheme is due to appear in the September 2003 issue of the IEEE Transactions on Circuits and Systems.

Long Term Goals

Our long term goal is to understand how to build complex and large-scale mixed analog-digital systems that compute efficiently. Hybrid analog-digital systems have the potential to be more efficient than either purely analog or purely digital systems. Some of our prior work has suggested that an underappreciated and important reason for the efficiency of the human brain, which only consumes 12W, is the hybrid and distributed nature of its architecture. Some of our other work has shown that the brain can be simultaneously analog and digital at the same time. How do we build biologically inspired artificial systems that approach the tremendous efficiencies in computation that biological systems have? For example, how does one build a high-precision digital/decision system using many moderate-precision analog parts, something that our brain does effortlessly and efficiently?

Scaling principles such as the distribution of information and information processing, signal restoration, and the creation of a complex computation via a sequence of simple ones are essential in building large-scale hybrid

systems. Spikes are natural for hybrid systems because spike time is analog while spike count is digital. They can naturally implement all of these scaling principles. We are interested in understanding how spikes may be exploited to build efficient hybrid computers.

A good concrete moderate-term goal that embodies our long-term goal is the creation of a sensory data processor: A signal processor that can perform high-precision computations on sensory data obtained from audition, vision, or olfaction through many moderate-precision analog operations. This task involves understanding efficient ways to perform analog signal restoration, efficient ways to do analog-to-digital conversion, efficient ways to build analog circuits that are immune to mixed-signal power-supply noise...and many other tasks that are the root of building mixed-signal systems that will scale to sizes of arbitrary complexity.

Best Accomplishments

Our best accomplishments this year involved: 1. The experimental realization of a time-based analog-to-digital converter with 10 bits of precision at 150uW of power consumption and a 100 kHz sampling frequency. These results represent more than a four fold improvement in the number of output levels of the converter from the past year and a significant improvement in speed and power consumption. They were due to algorithmic and circuit improvements in the converter. The final version of this converter, which we are awaiting from fabrication, promises to be the most power efficient converter ever reported in its speed and precision range. 2. The development of an analog VLSI medium-term analog memory cell for use in learning and adaptive analog systems. The memory cell consumes over three orders of magnitude less power than existing cells, while achieving 11-bits of sampled precision and 12-bit hold times of 45 seconds, three times longer than the nearest competitor. We expect to achieve 12-bits of sampled precision in our next iteration, making our work the world's leading medium-term analog memory cell implementation in all respects.

Scientific and Technical Objectives

Our scientific objective is to understand how to architect large-scale mixed-signal spiking neuronal systems in hardware in such a way that they will scale to systems of arbitrarily large complexity.

Our technical objective is to create an sensory data processor: A signal processor that can perform high-precision computations through many moderate-precision analog operations. Such processors have the potential to perform signal processing with a vast reduction in power. The overall task involves understanding efficient ways to perform analog signal restoration, efficient ways to do analog-to-digital conversion, efficient ways to build analog circuits that are immune to mixed-signal power-supply noise, efficient ways to use spiking neurons to build pattern recognition and learning circuits...and many other tasks that are the root of building mixed-signal systems that will scale to sizes of arbitrary complexity.

Approach Our approach is to first focus on building on various circuit building blocks that are important in the scaling of mixed-signal spiking systems: analog signal restoration blocks, ultra low power analog-to-digital converters, efficient comparators, power-supply immune circuits, pattern recognition circuits. Then we will focus on putting these blocks together to build systems in a general-purpose sensory data processor.
Impact/Applications There are 3 applications that we are actively pursuing:

1. Ultra low-power analog-to-digital conversion.
2. Analog Memory Elements for Learning Systems and Adaptive Analog Systems
3. An Ultra-Power-Efficient Asynchronous Comparator. This comparator is broadly useful in synchronous Rectifiers in power electronics, sensor interfaces, oscilloscope triggers, some types of analog-to-digital converters, and spiking-neuron circuits.

In addition, some of the circuits developed in this project are being explored for use in a bionic ear processor for the deaf.

Results

1. Capacitive-refresh analog storage techniques do not offer a power or area savings over traditional digital methods. However, simple leakage reduction or cancellation techniques can offer huge wins in terms of power and area over digital techniques in situations where long-term analog storage is not necessary (i.e. in circuits that adapt and learn from their environment in real time). Ultra-low power medium-term storage cells can be used as the basic storage element in fully-refreshed long-term storage cells, offering an energy savings to the overall system due to the lower refresh rate.

2. Adaptive mixed-signal systems with a handful of state variables are capable of amazingly complex and powerful behavior due to their nonlinear nature: The difference equations that govern our adaptive comparator may be represented by a unimodal discrete map. Therefore, near the limits of the scheme's stable regime of operation, we were able to experimentally confirm that there was a period-doubling route to chaos.

Technology Transfer

1. My low-power spike-based analog-to-digital converter is being explored for use in a low-power bionic ear for the deaf.

2. The spike-based analog-to-digital converter could possibly be the world's most energy-efficient converter in the near future or at least be very efficient.

3. Analog memory elements have use in several learning systems and adaptive analog systems that need to store analog parameters.

4. The adaptive comparator has applications in synchronous rectifier circuits in power electronic systems, in sensor interfaces, and in oscilloscope triggers.

Journals Articles

Sarpeshkar, R, and O'Halloran, M. (2002), Scalable Hybrid Computation with Spikes, Neural Computation, Vol. 14, No. 9, September 2002.

Mevay, A. and Sarpeshkar, R. (2003), Predictive Comparators with Adaptive Control, IEEE Transactions on Circuits and Systems II, Analog and Digital Signal Processing, in press, September 2003

Books or Chapters
None.

Technical Reports

None.

Patents Issued and Pending
None.

Honors/Awards
2003 MIT Junior Bose Award for Excellence in Teaching, 2001 Packard Award for Outstanding Young Faculty, 2001 ONR Young Investigator Award

References
None.

Related Projects
None.

****